

User Facilities as Interoperable Tools (for Science)

NUFO 2016

David Skinner

Strategic Partnerships Lead

National Energy Research Scientific Computing (NERSC)



Overview

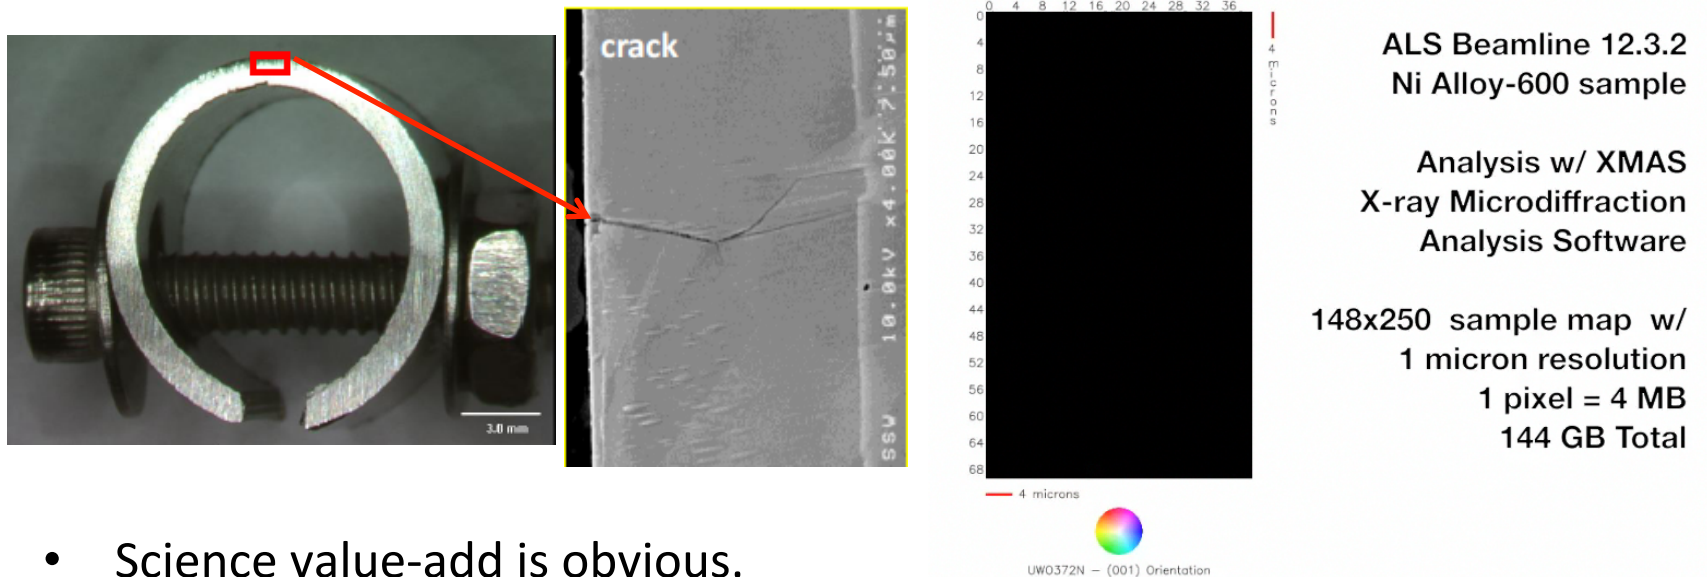
- National Energy Research Scientific Computing
- Facility founded in 1974, focused on open science, serves Office of Science supported research
- Fields 1-2 Large HPC systems at a time
- Great place to observe many science teams develop and run workflows involving simulation and data analysis.

Inter-facility Perspectives

- NERSC
 - HPC facility providing 3B+ hours to 5K+ users in ~700 projects enabling ~2K publications a year.
 - In transition to energy efficient architectures
 - Increasingly engaged in inter-facility science.
- ESnet
 - Built for inter-facility science. DTN, DMZ, SDN
 - Deploying capabilities to move multi-TB data sets.
 - Not at NUFO this year, but Lauren says “hi”
- ALS, JGI, and MolecularFoundry
 - Increasingly connected in ways we’ll describe.

Inter-facility Science Opportunities

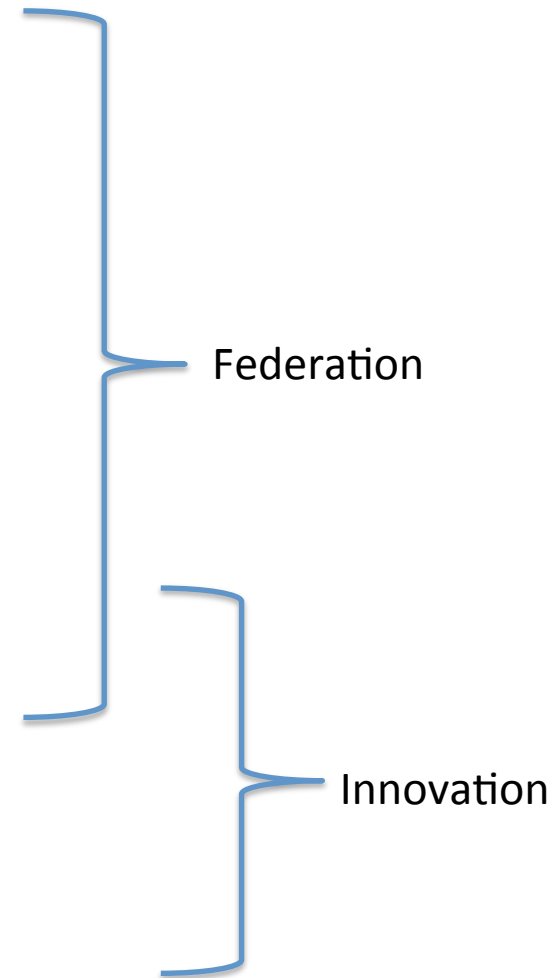
- There are many. I'll describe only a few.
- This particular one got my attention in 2012
 - Pilot study in connecting a beamline to HPC



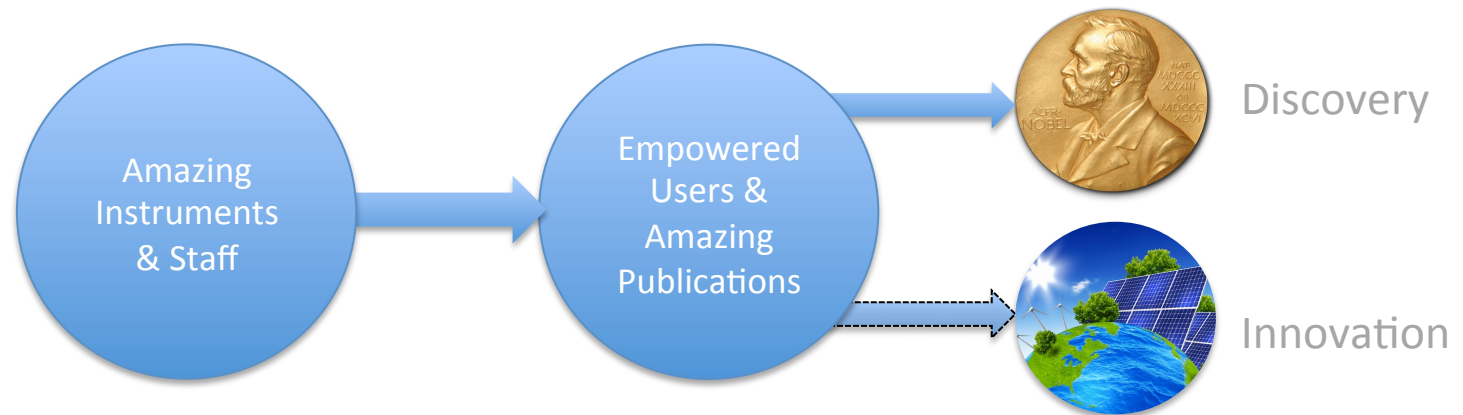
- Science value-add is obvious.
- Also many challenges: How to coordinate projects, accounts, allocations, and production workflow?
- Expanding operational modes of facilities is challenging work.

Inter-facility Challenges, a taxonomy

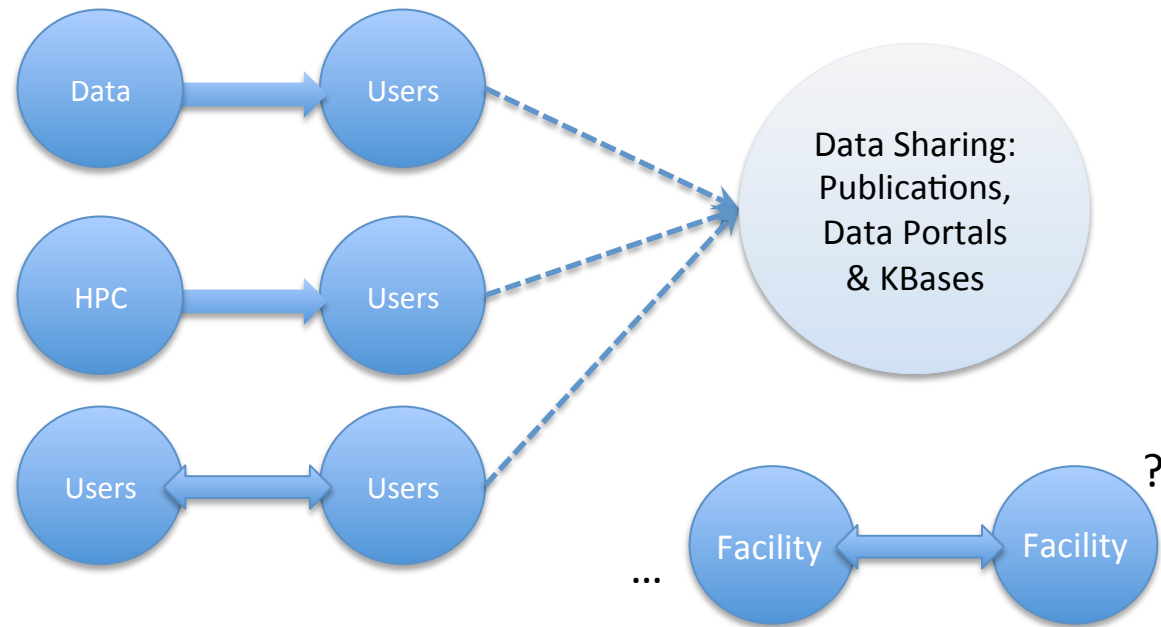
- Proposals
 - Multiple facilities means multiple proposals
- Agreements
 - Signed UA's at multiple facilities
- Accounts
 - User accounts, project accounts
 - UIDs and GIDs
- Scheduling
 - Overlap beam time w/ compute time
- Operational Modalities
 - IF may ask your to try new things
- Data management
 - Data pipelines, DMPs, and portals



What is a user facility? DOE SC Perspective



- Experiment & Observation
- Computing
- Networking



Inter-Facility Science

Opportunities and Challenges

Examples from ALS, JGI, Telescopes,
and LCLS

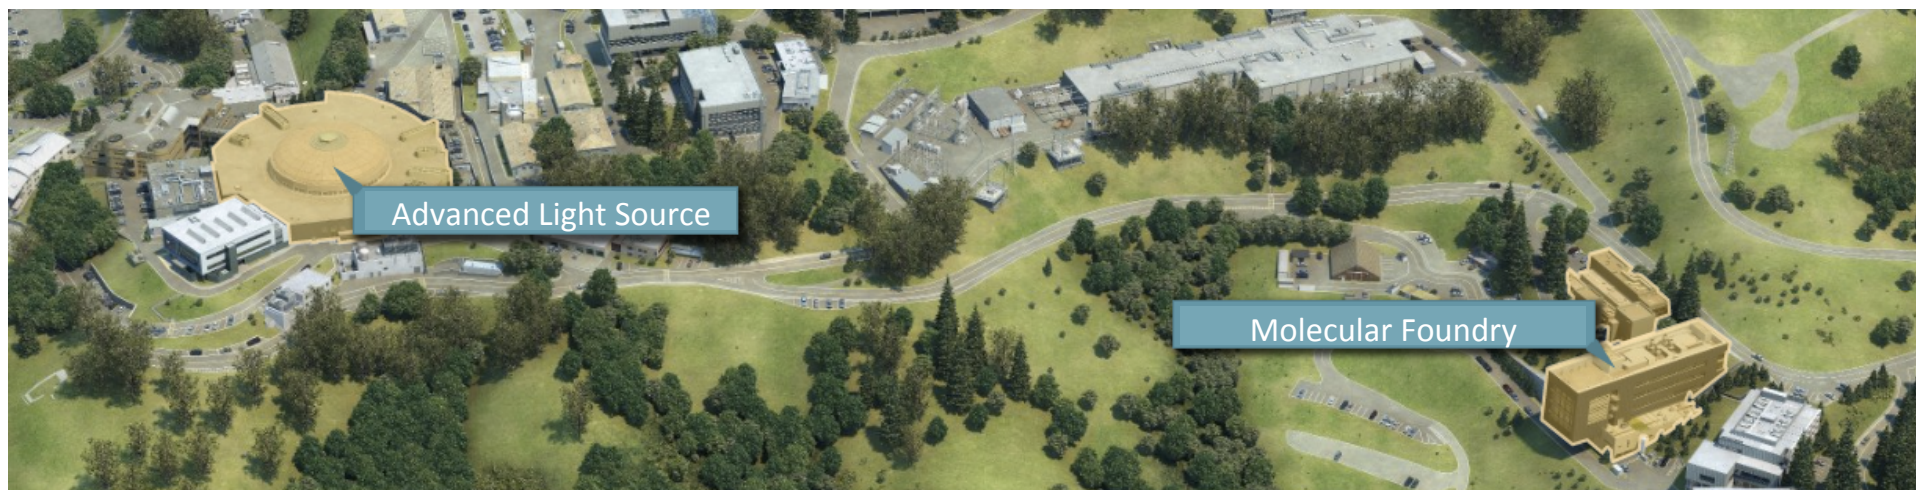
Molecular Foundry/Advanced Light Source

Coordinated Support for Collaborative User Access

Since 2010: MF-ALS MOU lets users request ALS beamtime in an MF proposal and vice versa.

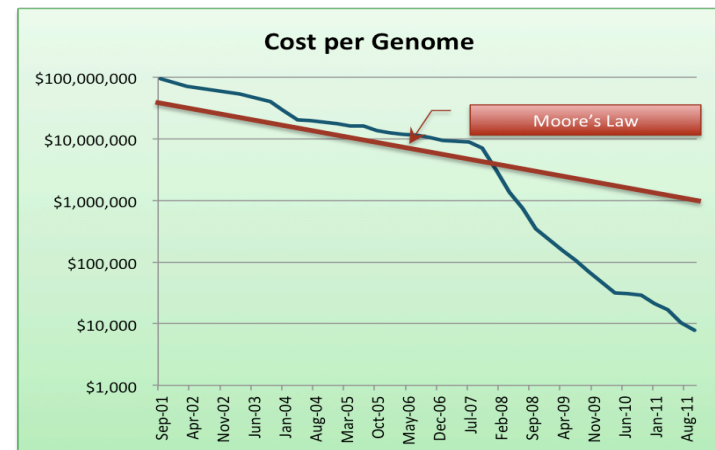
New in 2015/2016: Increased coordination between MF & ALS User Offices to facilitate user collaborative access

- MF & ALS user proposal databases connected
- Allows automatic transfer of user proposals between facilities
- Enables user to schedule training, beamtime, lab access, etc., without additional administrative overhead
- Ensures cross-facility access adheres to safety requirements



JGI's Transition to HPC

- NERSC and JGI formed partnership in 2010
- JGI was struggling to keep pace with the growing demand from sequencers.
- Federation: Merged accounts, UIDs, GIDs, software & staff
- A heavy lift but solid facility-facility federation



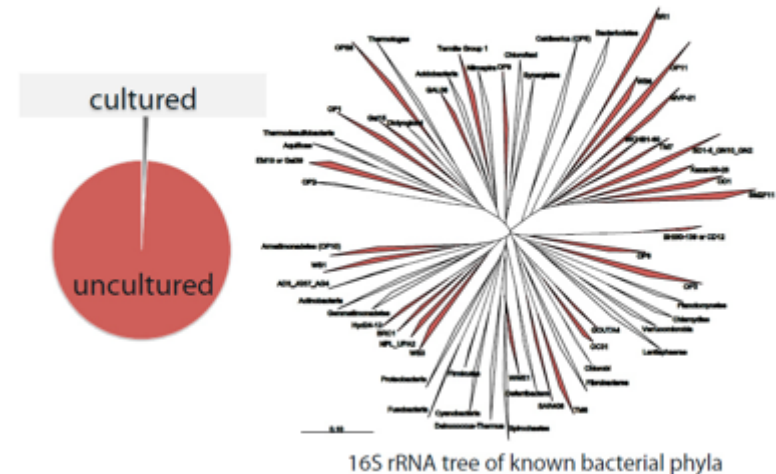
Source: National Human Genome Research Institute

NERSC supports the computing and data requirements of the Joint Genome Institute



- Since 2010 NERSC has supported the workflow and computing requirements of the JGI
- Sequencers generate up to 5 TBs per day
- Workflows consist of genome quality assurance, assembly, annotation, data archiving and curation
- JGI is increasing use of HPC systems.
- Application are benefiting from parallelism

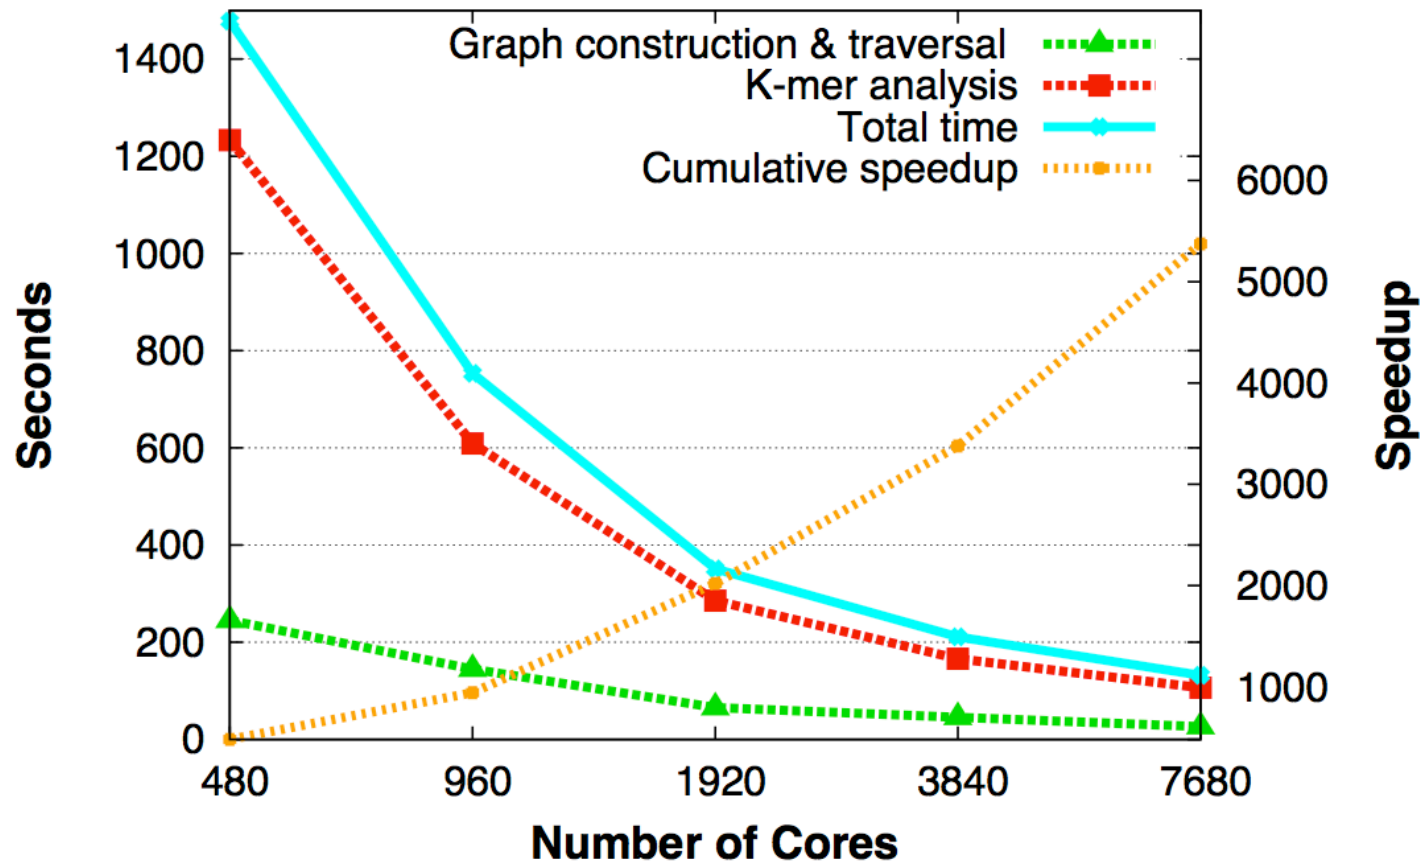
Our skewed view of the microbial world



phylogenetic tree reconstruction for the microbial dark matter project, *Nature* 2013



Genomics Software goes Exascale



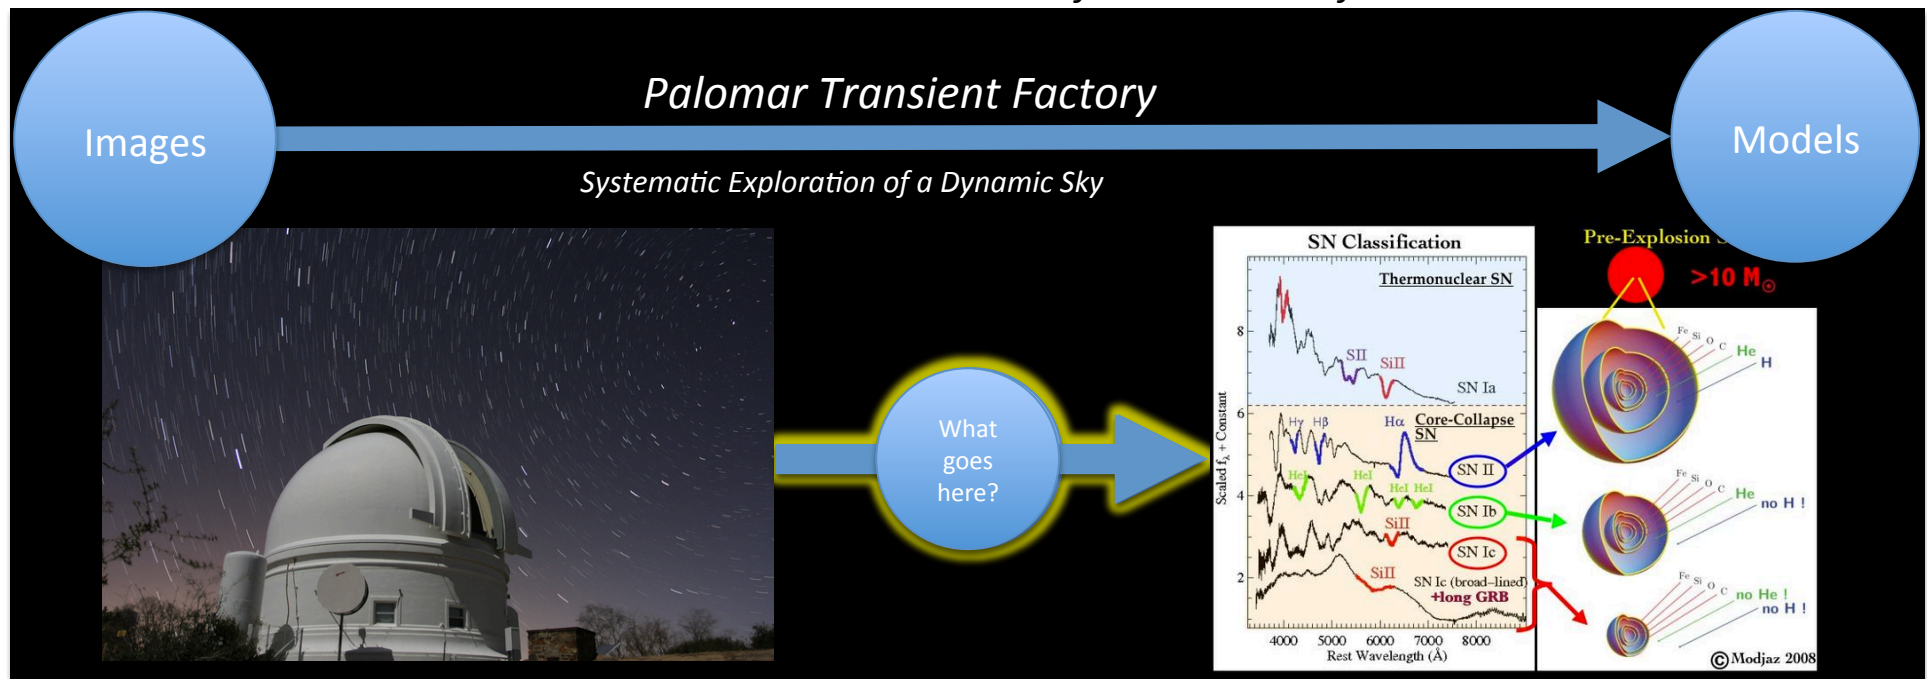
- UPC-based plant assembly code, the first scalable distributed assembler
- JGI and NERSC efforts move this to code to many-core architectures (KNL)

Data Analysis Above the Cloud(s)



The supernova of a generation 8/24/11

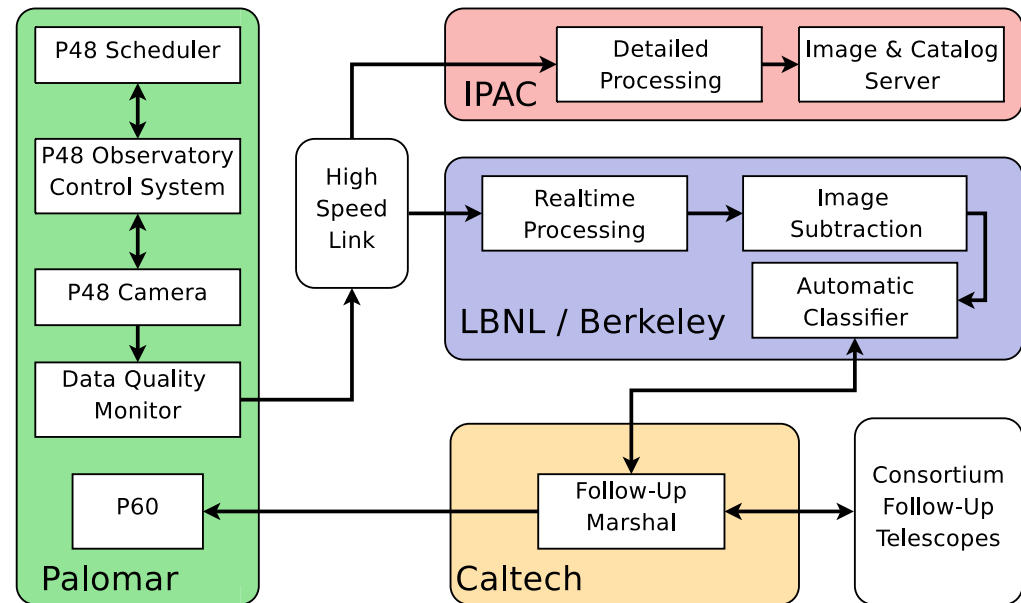
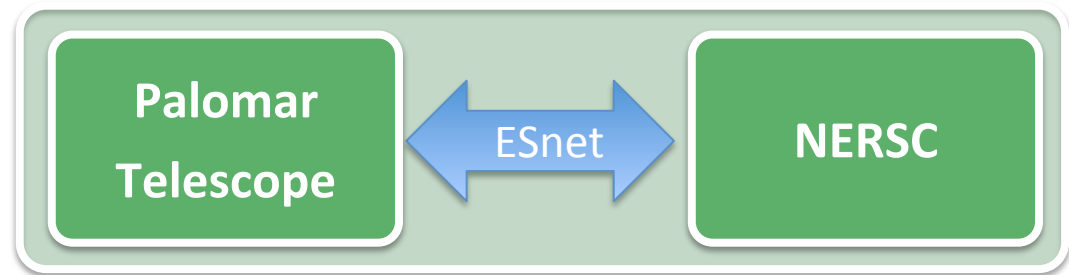
"The rapid discovery and classification of SN2011fe—all on the same night—is a testament to the great teamwork between all the researchers from over a half a dozen institutions"



The scientific legacy of PTF includes discovery of supernova 2011fe, the youngest Type Ia ever observed; the identification of new classes of super-luminous supernovae and Calcium-rich transients; discovery of a pre-explosion outburst from a Type IIn supernovae; discovery of many novae, asteroids, and variable stars.

Telescope Inter-facility Workflow

PTF Science : Key Projects	
Various SNe	Dwarf novae
Transients in nearby galaxies	Core collapse SNe
RR Lyrae	Solar system objects
CVs	AGN
AM CVn	Blazars
Galactic dynamics	LIGO & Neutrino transients
Flare stars	Hostless transients
Nearby star kinematics	Orphan GRB afterglows
Rotation in clusters	Eclipsing stars and planets
Tidal events	H-alpha sky-survey



50 GB/night

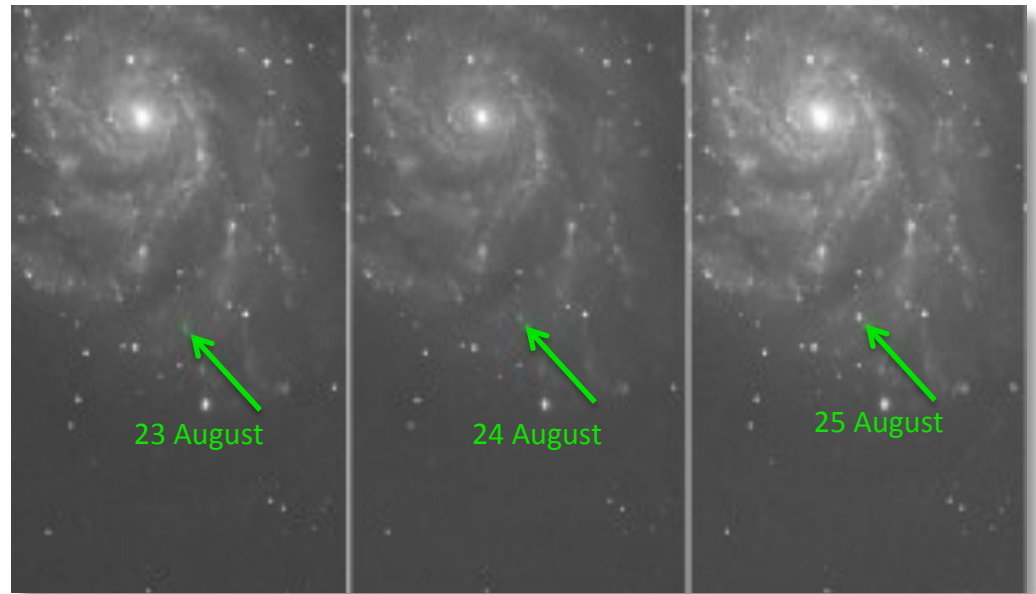
PTF: Fast Workflows → More Science

Observed within hours of its explosion, 11KLY will allow a rare glimpse at the supernova's outer layers that contain hints about what kind of star exploded

The last time a supernova of this type occurred so close was 1972. Before that: 1937, 1898 and 1572.

“The supernovae of a generation”

- Josh Bloom UCB



“Before” and “After” images showing the 11KLY discovery.

NERSC resources used:

NERSC data transfer nodes accept 300GB/night; HPSS archives;

Data (> 150 TB and growing) reside on NERSC Global Filesystem;

NERSC HPC runs subtraction pipeline to process images and detect new transients;

Science gateway nodes serve the DeepSky database to astronomers worldwide.

**Peter Nugent project PI and
Realtime Transient Detection Lead**

Light Source Data Analysis



Fast Feedback Analysis

- Get the data from the device
 - Quick pass/fail filtering, QA in a time-critical mode, data reduction, provenance
 - From detector to disk
- Experiment-driven
- Instrument operation
- Detector R&D agenda

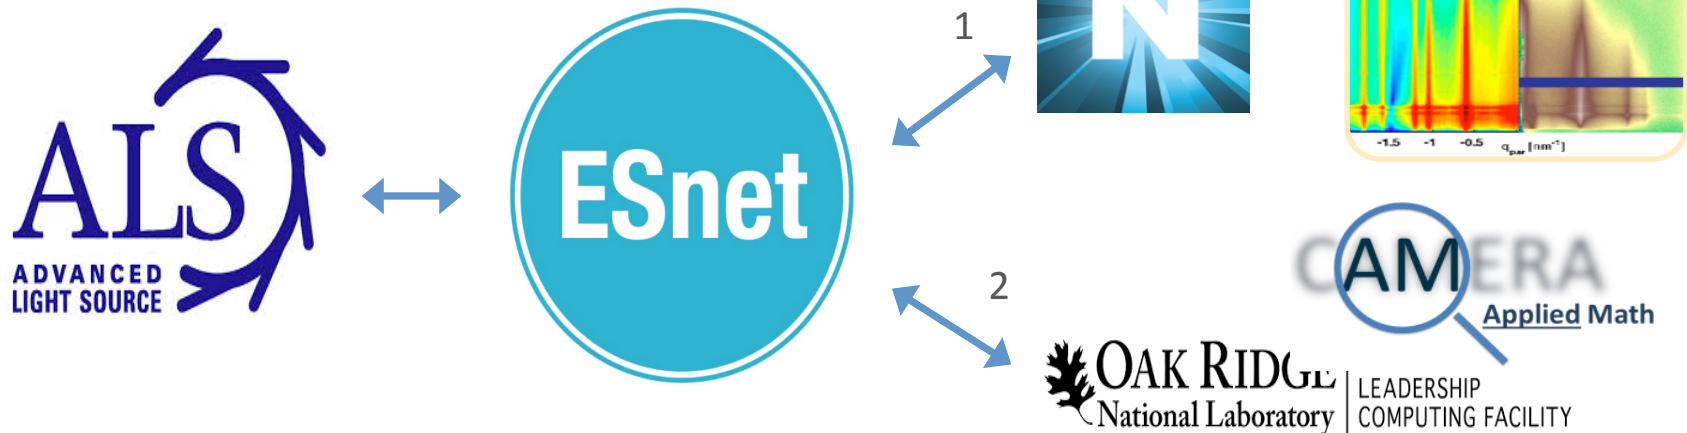
Scalable Offline Analysis

- Get the science from the data
 - Reconstructing and modeling from data. Time-to-solution is critical for large data sets
 - From disk to community
- Community-driven
- Data processing and Reuse
- Data analytics R&D agenda

Inter-facility X-ray Science

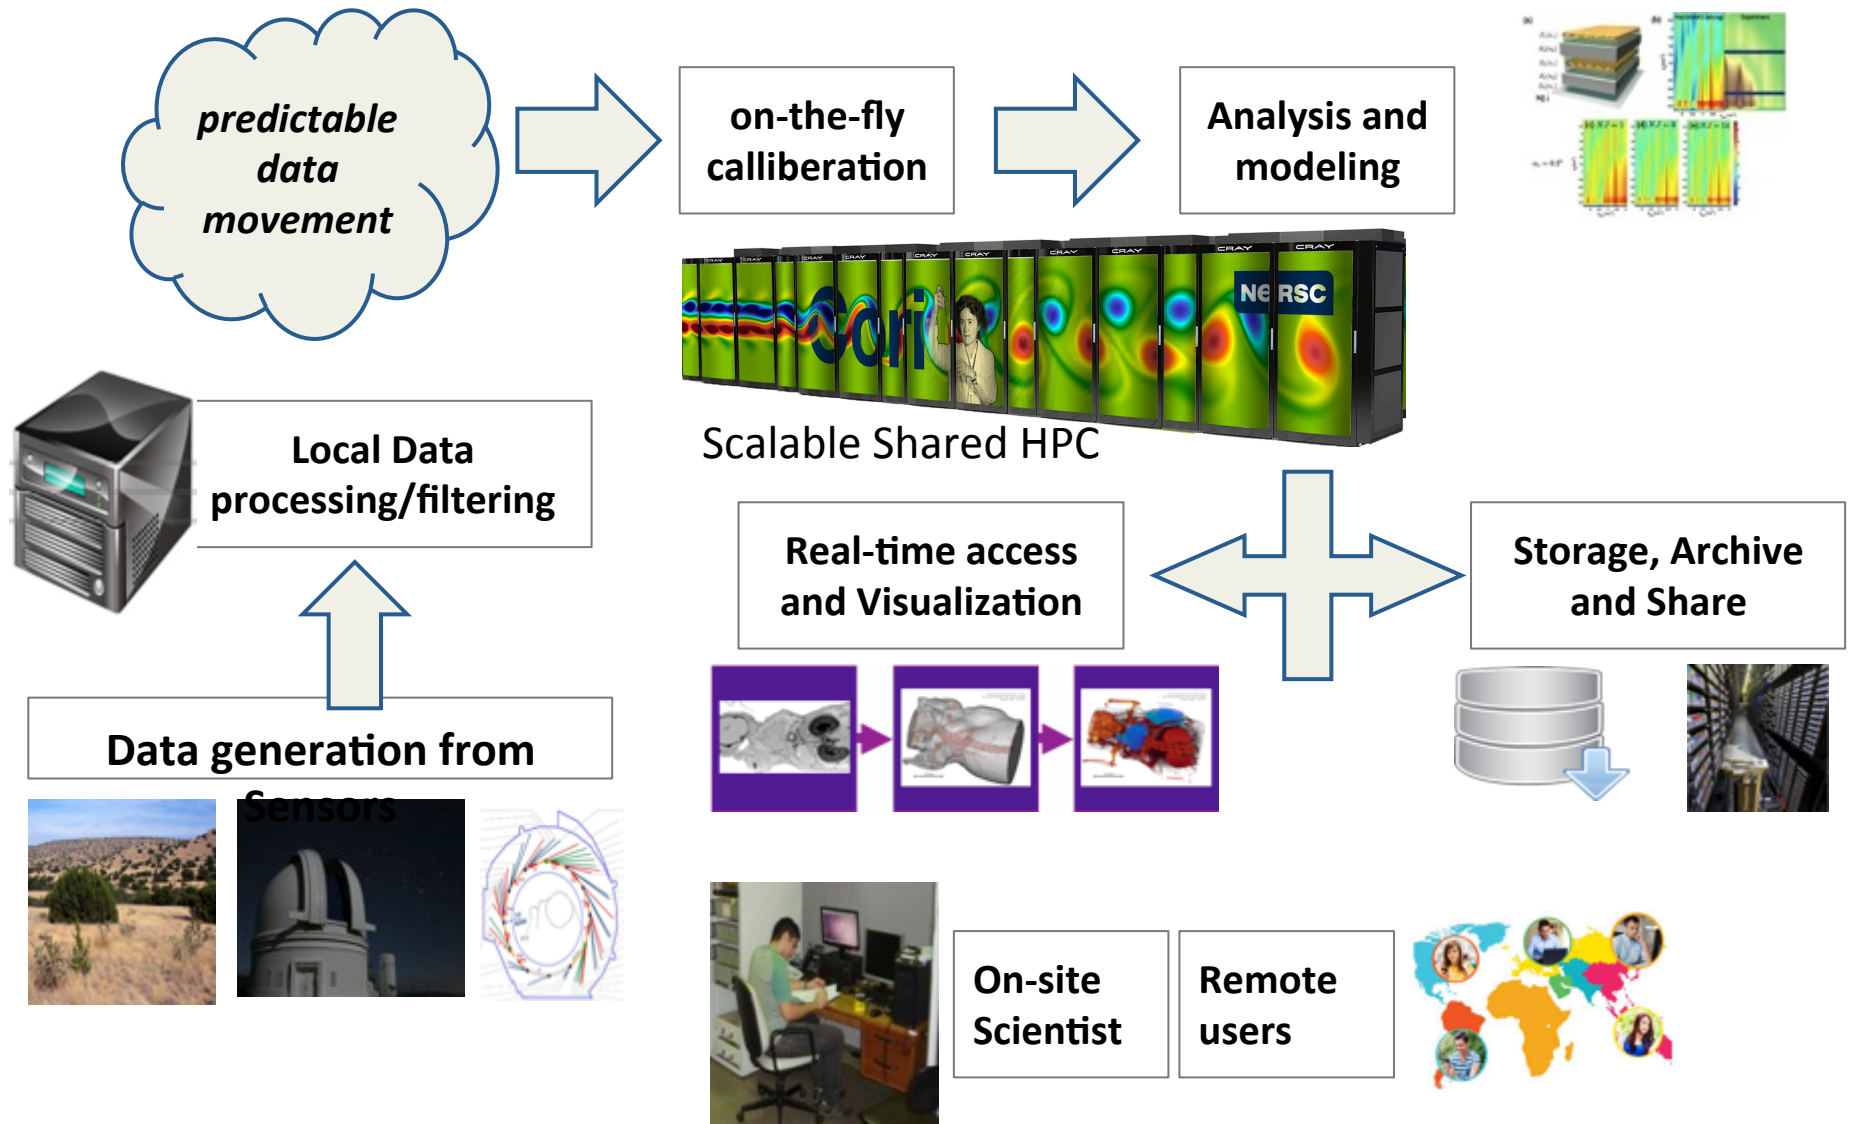
**‘Eliminate boundaries
between the Scientist and the
world’s best algorithms
running on the best
architecture for that code’**

**Real-time analysis of ‘slot-die’ technique for
printing organic photovoltaics, using ALS +
NERSC (SPOT Suite for reduction, remeshing,
analysis) + OLCF (HipGISAXS running on Titan
w/ 8000 GPUs).**

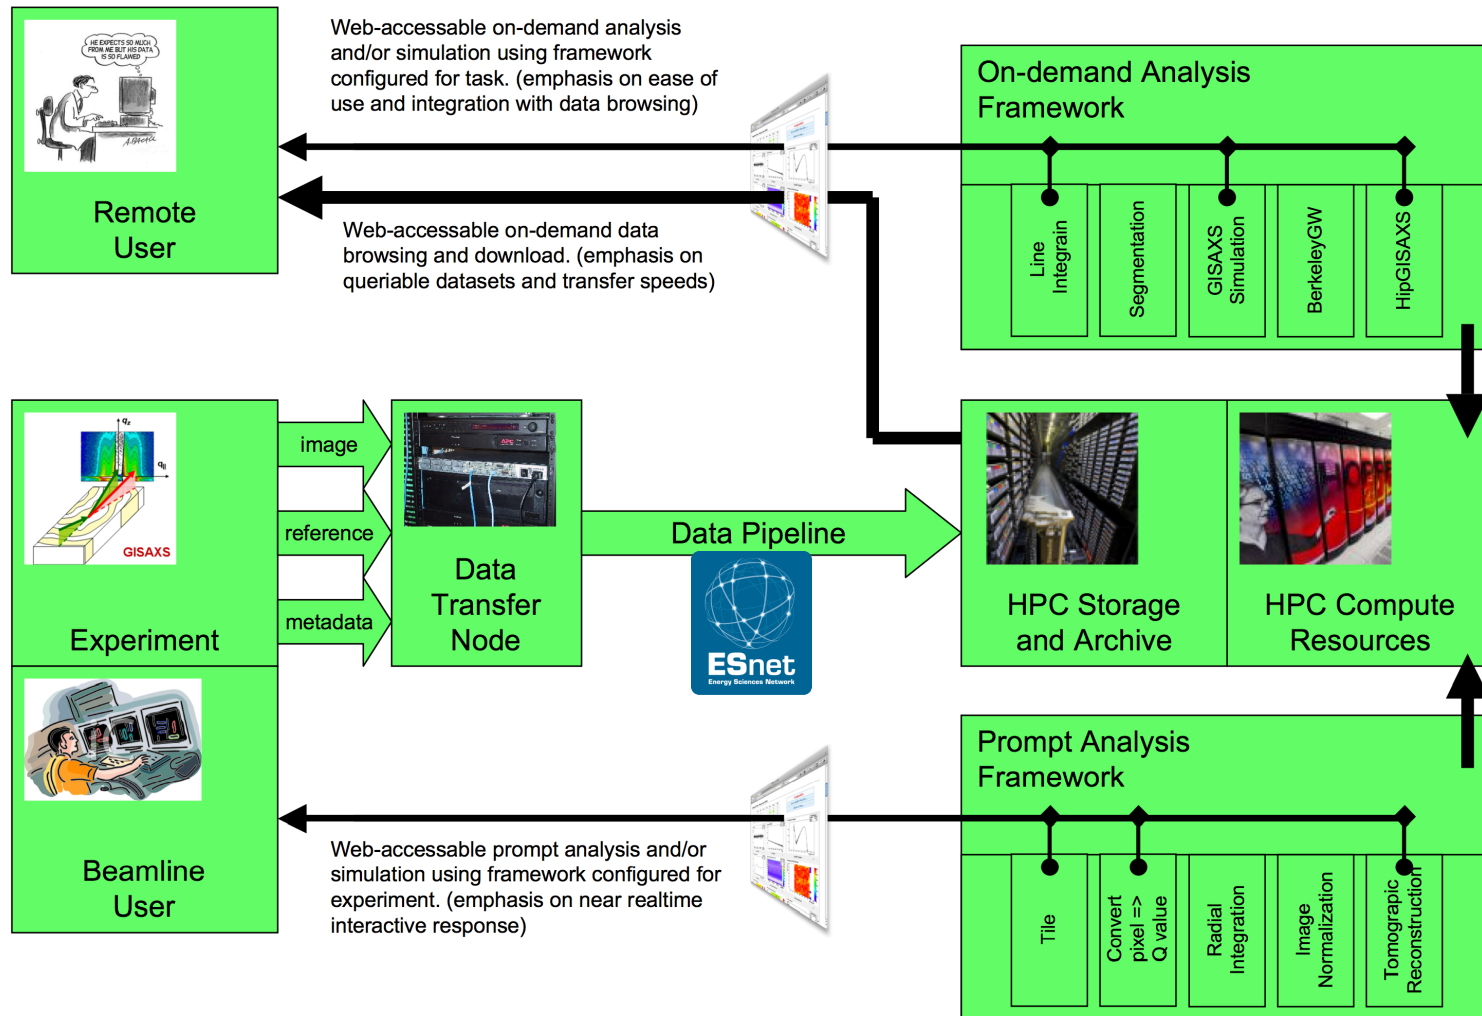


<http://www.es.net/news-and-publications/esnet-news/2015/esnet-paves-way-for-hpc-superfacility-real-time-beamline-experiments/> Results presented at March 2015 meeting of American Physical Society by Alex Hexemer. Additional DOE contributions: **GLOBUS** (ANL), **CAMERA** (Berkeley Lab)

Lightsource Workflow is Complex Test-case in Federation



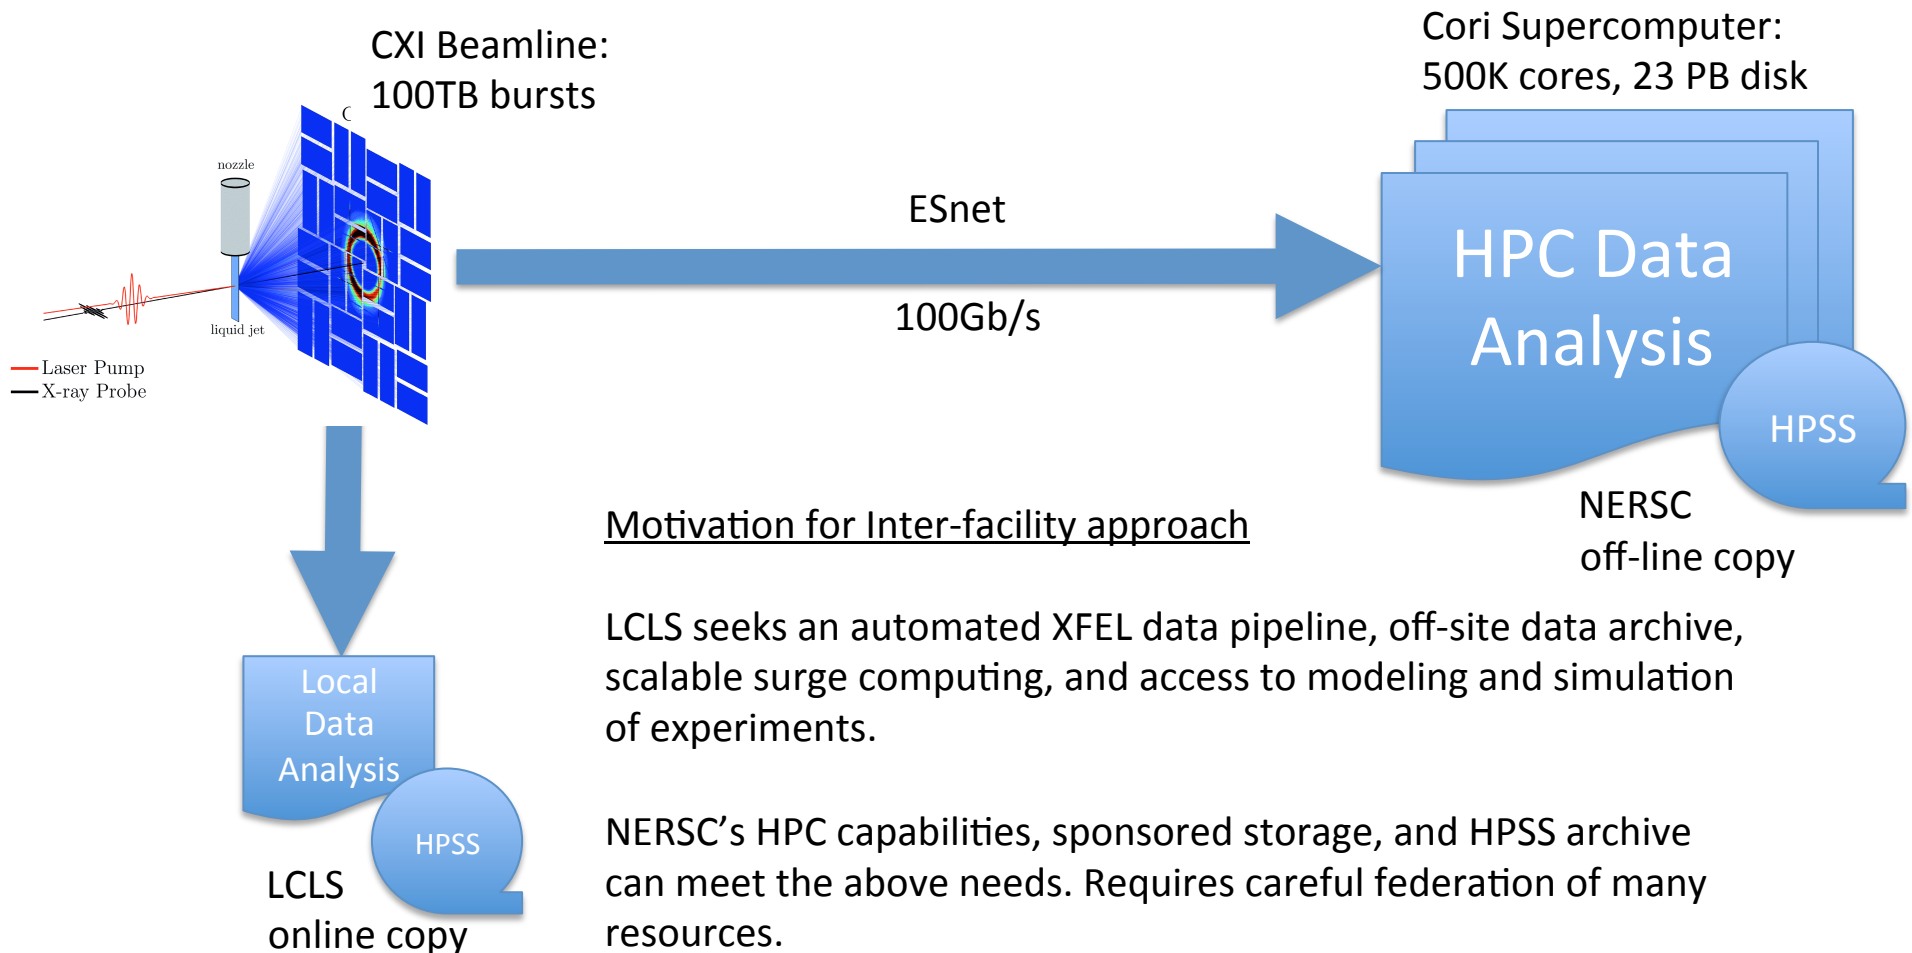
ALS Scientific Workflow Design



Craig Tull, Simone Patton, Jack Deslippe, Lavanya Ramakrishna, Alex Hexemer, Dula Parkinson, ...

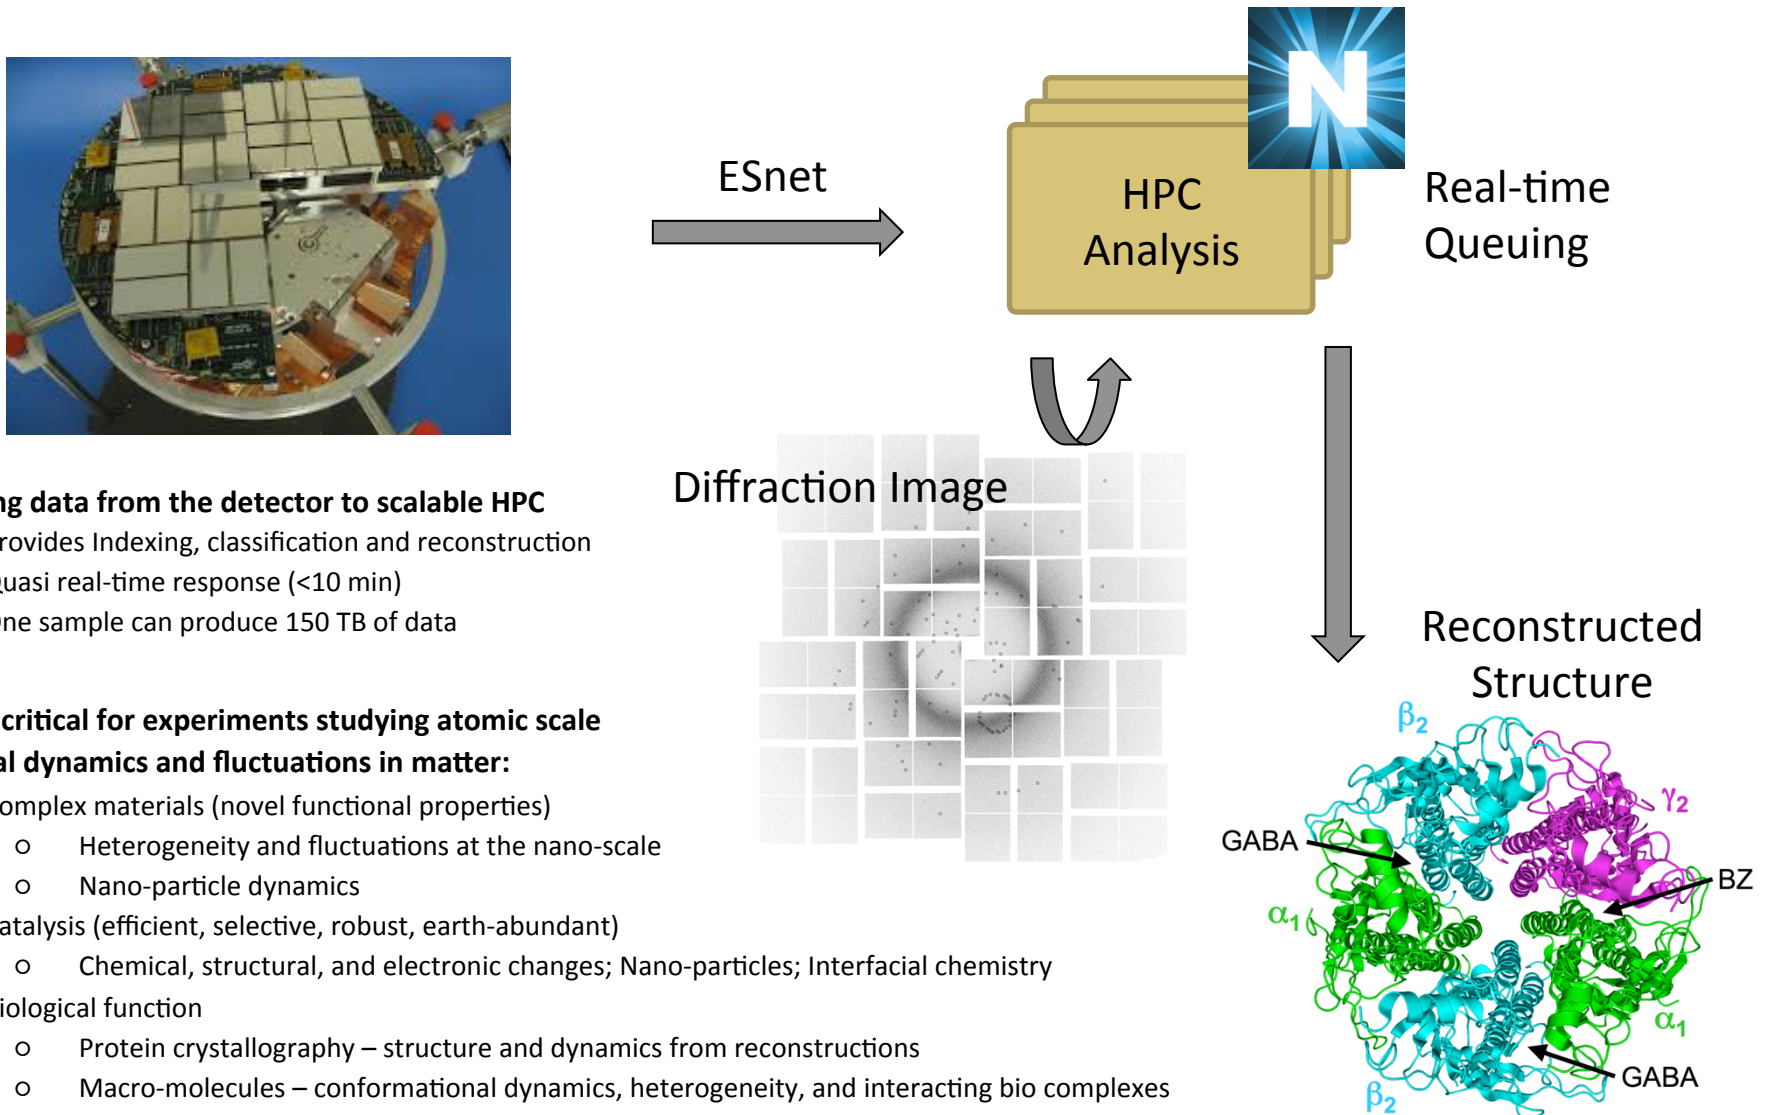
LCLS/NERSC Inter-facility XFEL Science

“Expand the impact of DOE’s experimental facilities by enabling new models of access, and discovery.”



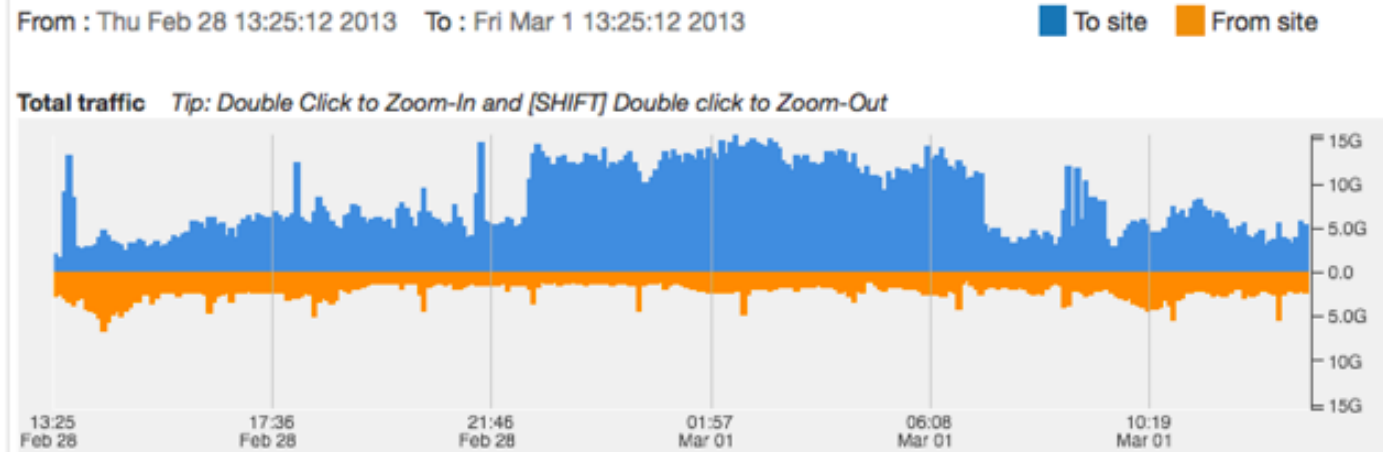
XFEL Usecase in Burst Processing

Providing atomic-scale vision to researchers at the beamline in < 10 min

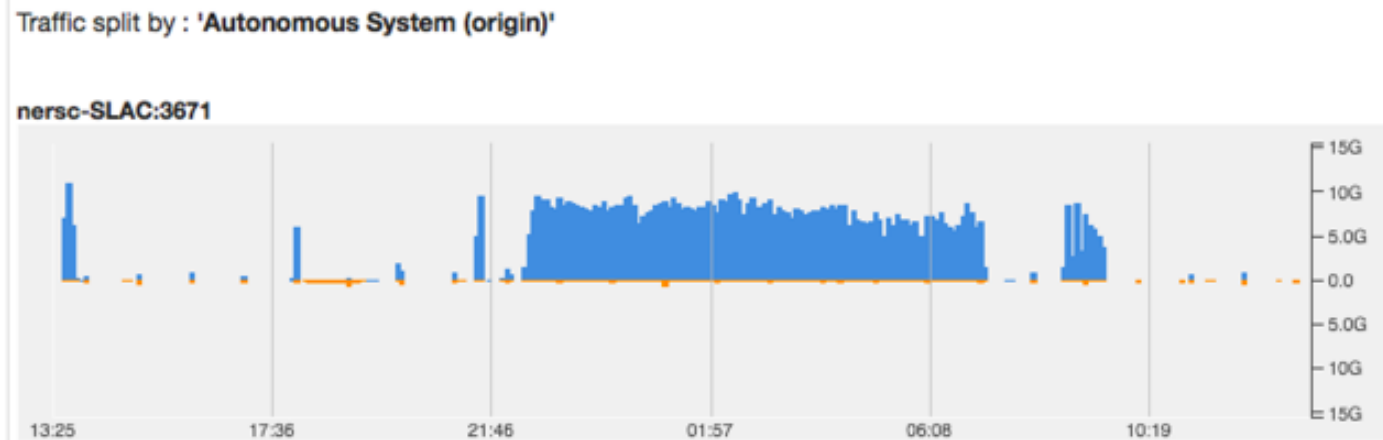


Cross Bay Data Transfer

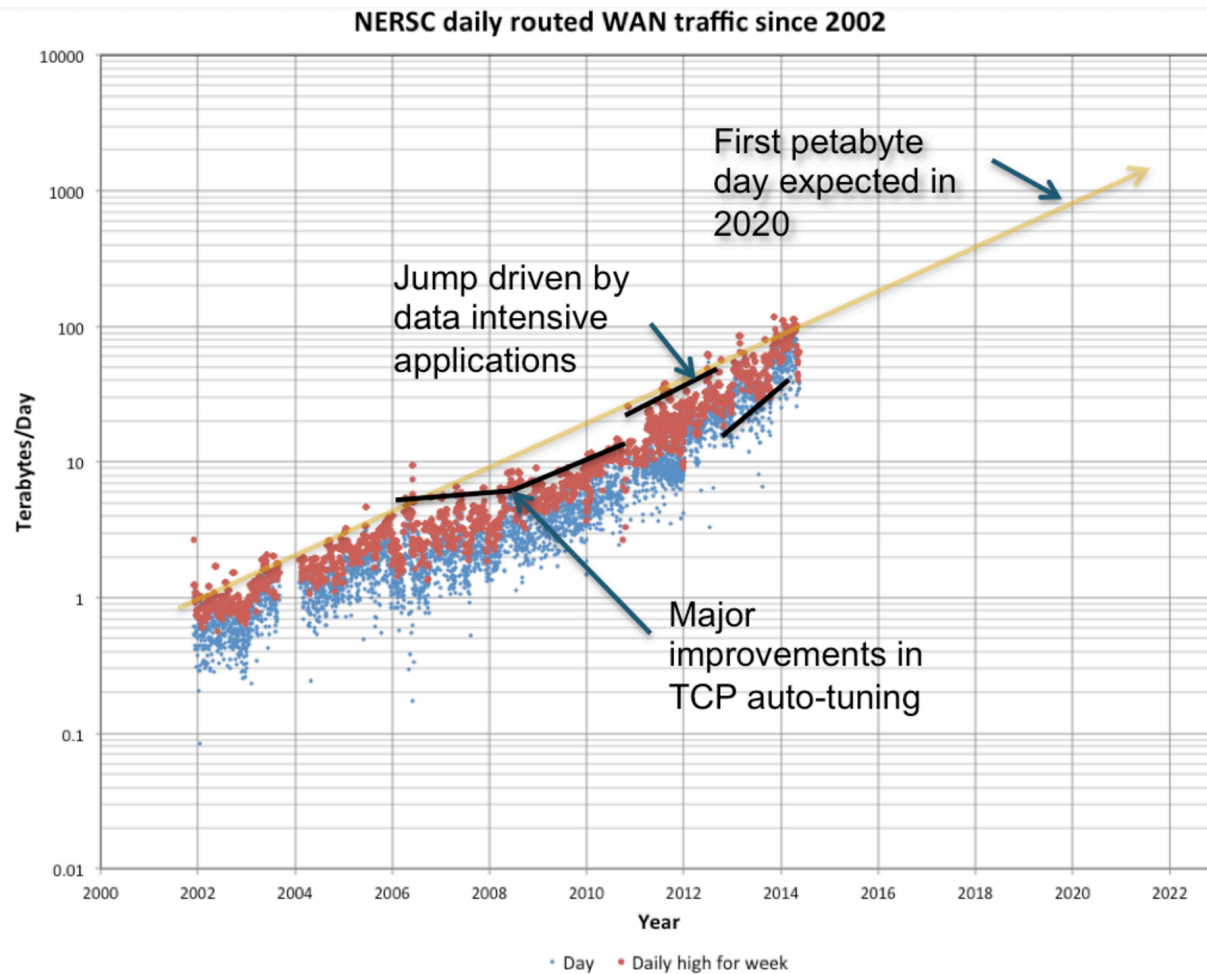
All NERSC
Traffic



Photosystem II
X-Ray Study



Yes, ESnet can move your data



Summary

Our Constellation of User Facilities is Increasingly Connected

World-class genomics, photon science, computing, nano-science and networking.

Increasingly facilities find user science agendas that require multiple facilities.

Need composed interoperation of science tools.



The Joint Genome Institute



The Advanced Light Source



NERSC



The Molecular Foundry



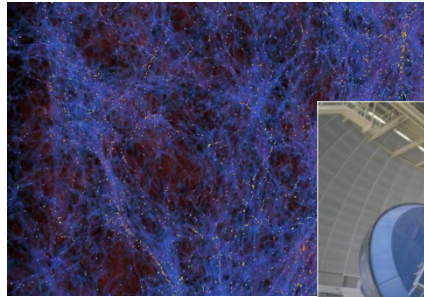
ESnet

Nearly 10,000 users of LBNL facilities

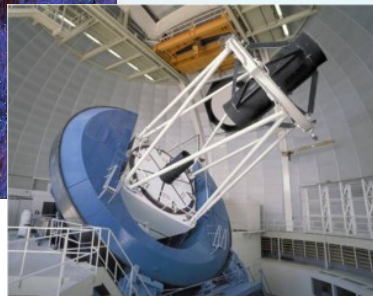
Many facilities adopting digital infrastructures

Good time to sort out federation issues

Federating experimental science with large scale data analysis & simulations will enable new discoveries

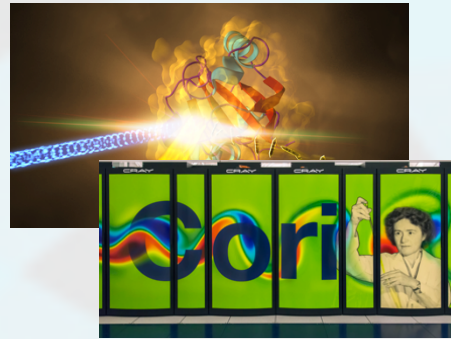


*Nyx simulation of
Lyman alpha
forest*



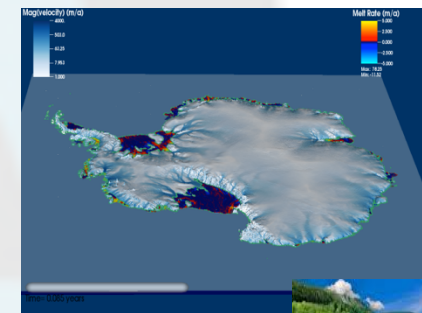
*Kitt Peak National
Observatory's Mayall 4-meter
telescope, planned site of the
DESI experiment*

XFEL photon science



Real-time Supercomputing

Climate modeling simulations



*Biology and
Sensor data*



*Esnet &
Science DMZ*

Crosscutting Concepts:

- User facilities (at their own pace) are becoming connected (HEP, JGI, ALS, LCLS..)
- Science DMZ model and an ecosystem of DTNs now reach user facilities at 100Gb/s.
- LBNL seeks to demonstrate discovery potential of seamless facility inter-operation
- Need to address gaps in automation and federation. NUFO best practices?